Supplementary: Neural Convolutional Surfaces

Luca Morreale[1*] Noam Aigerman[2] Paul Guerrero[2] Vladimir G. Kim[2]
Niloy J. Mitra[1,2]
[1]University College London [2]Adobe Research

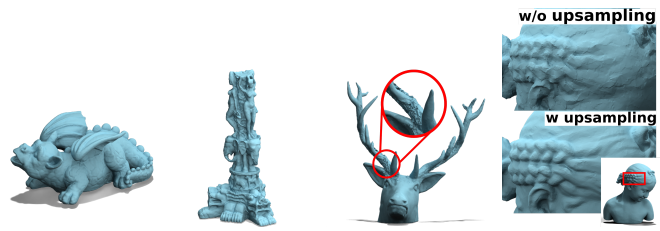

(a) Genus 2 **(b)** Genus 3 **(c)** ThinStructure **(d)** NoUpsampling

Figure 1. Results on high genus, thin structures, w/o upsampling.

1. Comparisons

Reconstruction comparison. For all baselines, [1–3], we used authors' implementations. See Figure 2 for qualitative a evaluation of our method, with (b), and without details (c).

Expressive power. Our construction is readily applicable to any genus, by cutting the mesh to a disk it is possible to reconstruct any surface. See Figure 1(a)(b) for reconstruction examples with different genus. However, Neural Convolutional Surfaces struggle to represent accurately thin structures. See Figure 1(c) for such a thin structure our framework is able to reproduce.

Finally, the upsampling is fundamental design choice for the CNN. Without upsampling the model is unable to capture details, see Figure 1(d).

2. Architecture Details

For the model f_ν, we used a 5-layer residual CNN with ReLU non-linearities. The fine MLP h_ξ uses a ReLU non-linearity after each layer except the last, and the coarse MLP g^c_ϕ uses Softplus activations. Please refer to Table 1 for complete architecture details of each model.

[*]Partially worked on the project during internship at Adobe Research

Table 1. Architecture details used for each shape presented in the paper.

	Coarse MLP g^c_ϕ	Per-patch Code Ω_i	CNN f_ν channels	Fine MLP h_ξ
Armadillo-100K	128-64-64	$8 \times 4 \times 4$	8	16-16
Bimba-100K	128-64	$8 \times 4 \times 4$	8	16-16
Dino-100K	64-64	$8 \times 8 \times 8$	8	16-16
Dragon-100K	128-64-64	$6 \times 6 \times 6$	8	16-16
Gargoyle-100K	64-64-64-64	$6 \times 4 \times 4$	6	16-16
Grog-100K	128-64-64	$8 \times 4 \times 4$	8	16-16
Seahorse-100K	128-64-64	$8 \times 8 \times 8$	8	16-16
Elephant-100K	128-64-64	$8 \times 6 \times 6$	8	16-16
Armadillo-1M	128-64-64	$64 \times 4 \times 4$	64	16-16
Bimba-1M	128-64	$64 \times 4 \times 4$	64	16-16
Dino-1M	64-64	$64 \times 8 \times 8$	64	16-16
Dragon-1M	128-64-64	$66 \times 6 \times 6$	64	16-16

References

[1] J. N. Martel, D. B. Lindell, C. Z. Lin, E. R. Chan, M. Monteiro, and G. Wetzstein. Acorn: Adaptive coordinate networks for neural scene representation. *arXiv preprint arXiv:2105.02788*, 2021. 1

[2] T. Takikawa, J. Litalien, K. Yin, K. Kreis, C. Loop, D. Nowrouzezahrai, A. Jacobson, M. McGuire, and S. Fidler. Neural geometric level of detail: Real-time rendering with implicit 3d shapes. In *Proc. CVPR*, pages 11358–11367, 2021. 1

[3] W. Yifan, L. Rahmann, and O. Sorkine-Hornung. Geometry-consistent neural shape representation with implicit displacement fields, 2021. 1

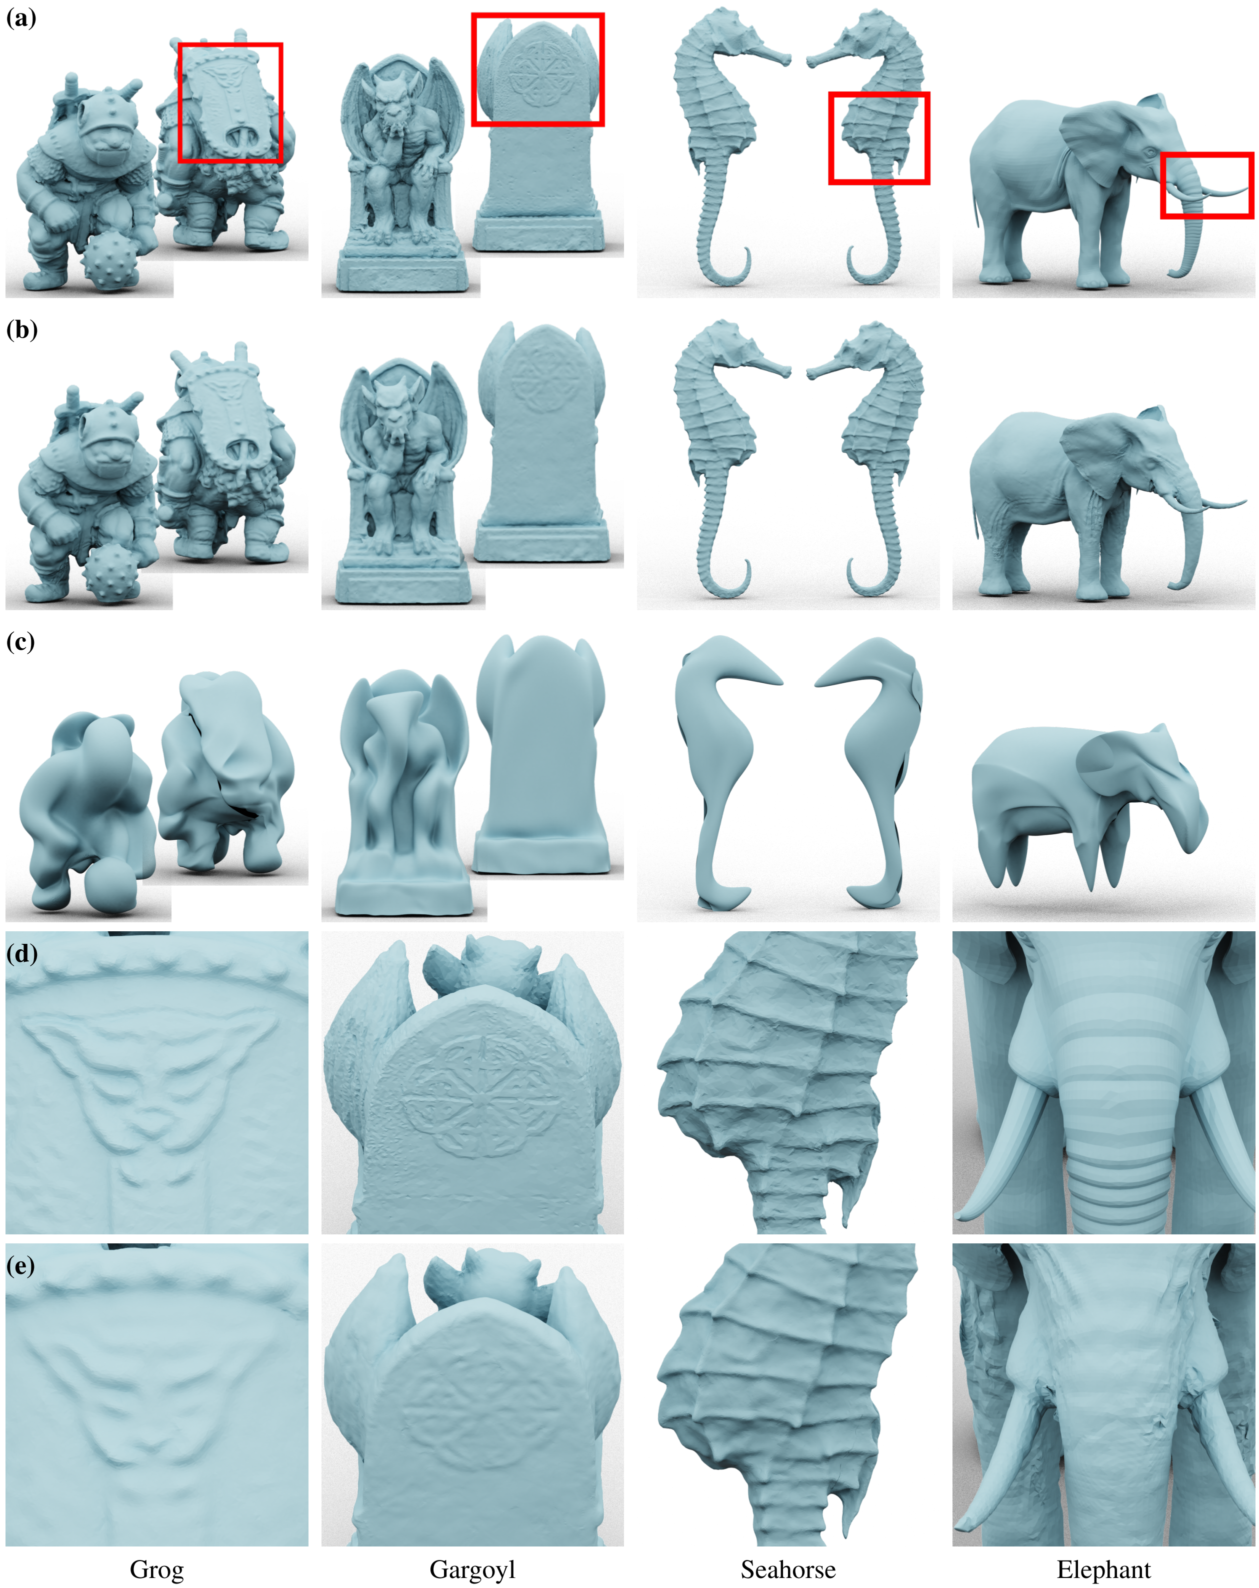

Figure 2. Representation quality -with (b) and without (c) details- of our method, compared with the ground truth model (a). We limit Neural Convolutional Surfaces to 100K parameters. We show inset zooms (e) of our reconstruction for further assessment, with corresponding inset zooms (d) for ground truth.